

Decision Accessibility and Structural Inaccessibility under Representational Collapse

Pierernesto Di Paolo - Independent Researcher - ORCID: 0009-0003-0149-4321

2026

Working Paper / Preprint for Zenodo Deposit

Version: 1.0

Document type: Theoretical working paper with operational audit metric

Recommended citation before DOI assignment:

Di Paolo, P. (2026). *Decision Accessibility and Structural Inaccessibility under Representational Collapse*. Working paper / preprint.

Recommended citation after DOI assignment:

Di Paolo, P. (2026). *Decision Accessibility and Structural Inaccessibility under Representational Collapse*. Zenodo. DOI to be inserted after publication.

Author Note

This working paper presents a formal-operational framework for identifying decision failures caused by representational collapse. It is designed as an autonomous preprint: it does not require prior reading of any related book, manuscript or repository. The work is theoretical in structure and audit-oriented in application. No clinical, legal, financial or technical advice is provided or implied.

Abstract

Decision systems do not operate directly on the full state of reality. They operate on representations: measurements, records, categories, evidentiary files, feature vectors, embeddings, scores, or institutional descriptions. This paper formalizes a class of decision failures that arise when a representation collapses distinct real states into the same operational configuration. The central claim is deliberately limited: if different real states become indistinguishable under the available representation while requiring mutually incompatible correct decisions, then no deterministic decision rule based only on that representation can guarantee correctness across all those states. This is not a claim about negligence, irrationality, bias, or poor optimization. It is a claim about the accessibility of correct decisions under a given representation.

The framework is defined by a state space X , a representation space Z , a representation map $C: X \rightarrow Z$, a decision set D , an availability correspondence $A^i(z) \subseteq D$, and a correctness correspondence $B(x) \subseteq D$. The paper separates three conditions that are often conflated: avoidable error, fiber-level structural ambiguity, and state-level necessary failure. The main theorem gives necessary and sufficient conditions for guaranteed correctness on a representational fiber. A decision rule can be guaranteed correct on all states represented by

the same z if and only if there exists at least one decision that is both available under z and correct for every state in the fiber $C^{-1}(z)$. When this common intersection is empty, correctness cannot be guaranteed without changing the representation, the available decision set, or the information architecture.

The framework is positioned against bounded rationality, information theory, cybernetics, and statistical decision theory. Unlike Simon's bounded rationality, the primary limitation is not located in the cognitive capacity of the agent, but in the discriminative structure of the representation. Unlike Shannon's communication theory, the focus is not information transmission in general, but decision-relevant information loss. Unlike standard accuracy-based evaluation, the framework identifies local representational regions where errors may be irreducible under the given representation. The paper introduces the Representational Collapse Index (RCI) as an audit metric for finite datasets or case archives. RCI estimates the proportion of cases located in representational classes containing incompatible correct decisions. The result is a formal-operational framework, not a universal theory of correctness across all domains.

Keywords: decision theory; representation; structural inaccessibility; representational collapse; information loss; bounded rationality; cybernetics; audit; artificial intelligence; legal decision-making; clinical decision-making.

1. Introduction

A decision system rarely decides from reality itself. It decides from a representation of reality. In medicine, the decision is based on symptoms, measurements, images, laboratory values and clinical records. In law, it is based on admissible evidence, procedural acts, pleadings, files and recognized facts. In artificial intelligence, it is based on features, embeddings, tokens, labels, sensor streams or learned internal states.

This creates a structural problem. A system may fail not because the decision-maker is irrational, not because the model is weak, and not because the decision rule is poorly optimized. It may fail because the representation does not preserve the distinction required for a correct decision.

The guiding question of this paper is: when does a correct decision become inaccessible because of the representation on which the system is forced to operate?

The answer is not that every loss of information is fatal. Most representations discard information. Some discarded information is irrelevant to the decision task. The critical case arises when the representation collapses states that require incompatible correct decisions. Then any rule that sees only the collapsed representation must treat those states identically, even though correctness requires treating them differently.

The contribution of this paper is fourfold. First, it gives a formal vocabulary for distinguishing avoidable error from structural inaccessibility. Second, it proves a fiber-level criterion for guaranteed correctness under representation constraints. Third, it introduces an operational audit metric, the Representational Collapse Index, for detecting decision-relevant collapse in finite samples. Fourth, it positions the result with respect to bounded rationality, information theory, cybernetics and decision theory.

The framework is intentionally limited. It does not define correctness universally. It does not claim that all errors are structural. It does not eliminate domain judgment. Its purpose is narrower: to identify cases in which no improvement of the decision rule alone can recover a distinction already absent from the representation.

2. Theoretical Positioning

2.1 Bounded rationality

Simon's bounded rationality identifies limits in the cognitive, informational and organizational capacities of the decision-maker. The bounded agent cannot search all alternatives, compute all consequences, or optimize under complete knowledge. The result is satisficing rather than perfect optimization (Simon, 1957).

The present framework shifts the locus of limitation. The primary constraint is not the agent's cognitive capacity, but the representation available to the agent or system. Even a fully rational agent cannot select a guaranteed correct decision when the representation collapses states requiring mutually incompatible correct decisions.

The distinction can be stated sharply: in bounded rationality, the decision-maker is limited; in structural inaccessibility, the representation is insufficiently discriminative for the task. This distinction matters operationally. If the problem is bounded rationality, the intervention may be better reasoning, better computation, better training or better incentives. If the problem is representational collapse, those interventions may fail unless the representation itself is changed.

2.2 Shannon and decision-relevant information loss

Shannon's theory of communication formalizes information, entropy and channel capacity (Shannon, 1948). In that framework, a channel may transmit, distort or lose distinctions. The present paper borrows the intuition of information loss but narrows it to decision accessibility.

A representation C may be non-injective without causing decision failure. If two states collapse into the same representation but have at least one common correct decision, the collapse may be harmless for that task. Collapse becomes structurally relevant only when the lost distinction separates incompatible correctness sets.

The relevant question is therefore not simply whether information has been lost, but whether information necessary for access to a correct decision has been lost. This is a decision-theoretic refinement of information loss.

2.3 Ashby and requisite variety

Ashby's law of requisite variety states that effective regulation requires sufficient variety in the regulator to absorb or counteract the variety of disturbances (Ashby, 1956). The present framework is compatible with this idea but focuses on a specific mechanism: insufficient representational variety.

A system may have many possible actions and still fail if its representation does not distinguish the states for which those actions are appropriate. Conversely, a representation may preserve many distinctions but still fail if the available decision set is too narrow.

Thus structural inaccessibility may arise from insufficient representational variety, insufficient decision variety, mismatch between representation and decision availability, or loss of the specific distinction required by the correctness criterion.

2.4 Statistical decision theory

Statistical decision theory studies decisions under uncertainty by combining actions, states, loss functions and probability distributions (Wald, 1950). The present framework is compatible with that tradition but emphasizes a different boundary.

Statistical decision theory can identify an optimal action under a posterior distribution. However, optimality in expectation does not imply guaranteed correctness for every real state compatible with the representation. If a representation z is compatible with states requiring incompatible correct actions, the best expected decision may still be structurally unable to be correct for all states in the fiber.

2.5 Arrow and structural impossibility

Arrow's impossibility theorem shows that certain desirable conditions cannot be jointly satisfied in collective preference aggregation (Arrow, 1951). The present framework is not a social choice theorem. It does not concern preference aggregation. The connection is methodological: both approaches identify failures that are structural rather than accidental.

Here the structure is not a voting rule. It is the relation between real states, representations, available decisions and correctness.

3. Formal Framework

Let:

- X be a non-empty set of real states;
- Z be a non-empty set of representations;
- D be a non-empty set of possible decisions;
- $C: X \rightarrow Z$ be a representation map;
- $A^i: Z \rightarrow P(D)$ be an availability correspondence;
- $B: X \rightarrow P(D)$ be a correctness correspondence.

For each state $x \in X$, the system does not directly observe x . It operates on $z = C(x)$.

The set $A^i(z)$ contains the decisions available under representation z . Availability may be physical, legal, institutional, computational, procedural or epistemic. The set $B(x)$ contains the decisions correct for state x . Correctness is domain-dependent. The framework does not assume that $B(x)$ is a singleton.

A deterministic decision rule is a function $f : Z \rightarrow D$. A rule is availability-respecting if $f(z) \in A^i(z)$ for every relevant z . The decision selected in state x is $f(C(x))$. A decision is correct at x if $f(C(x)) \in B(x)$.

4. Basic Definitions

Definition 1 - Representational equivalence. Two states $x_1, x_2 \in X$ are representationally equivalent under C if $C(x_1) = C(x_2)$. We write $x_1 \sim_C x_2$. This relation partitions X into representational fibers.

Definition 2 - Fiber. For $z \in Z$, the fiber generated by z is $F_z = C^{-1}(z) = \{x \in X : C(x) = z\}$. A fiber is the set of all real states that are indistinguishable under representation z .

Definition 3 - Representational collapse. A representational collapse occurs at z if $|F_z| \geq 2$. Equivalently, there exist $x_1 \neq x_2$ such that $C(x_1) = C(x_2) = z$. Representational collapse is not automatically pathological. It becomes decisionally relevant only when the collapsed states require incompatible decisions.

Definition 4 - Common correctness set of a fiber. For a fiber F_z , define its common correctness set as $K(z) = \bigcap_{x \in F_z} B(x)$. This is the set of decisions that are correct for every real state compatible with representation z .

Definition 5 - Available common correctness set. For a representation z , define $G(z) = A^i(z) \cap K(z) = A^i(z) \cap \bigcap_{x \in F_z} B(x)$. This is the set of decisions that are both available under z and correct for every state in the fiber.

Definition 6 - Decision conflict. Two states $x_1, x_2 \in X$ are in decision conflict if $B(x_1) \cap B(x_2) = \emptyset$. No single decision is correct for both states.

Definition 7 - Decisionally incoherent fiber. A fiber F_z is decisionally incoherent if $K(z) = \emptyset$. Pairwise conflict is sufficient for incoherence but not necessary in all set systems when more than two states are involved.

Definition 8 - Avoidable error. A system commits an avoidable error at x when the selected decision is incorrect, although at least one correct available decision exists: $f(C(x)) \notin B(x)$ and $A^i(C(x)) \cap B(x) \neq \emptyset$.

Definition 9 - State-level necessary failure. A system is in state-level necessary failure at x when no decision available under the representation is correct for that state: $A^i(C(x)) \cap B(x) = \emptyset$. This is a definition, not a theorem.

Definition 10 - Fiber-level structural ambiguity. A representation z has fiber-level structural ambiguity when $G(z) = \emptyset$. This means that no single available decision is guaranteed correct for all real states compatible with z .

Definition 11 - Indeterminate case. A case is indeterminate when the relevant intersection cannot be established from available information. This may occur because $B(x)$ is unknown, disputed, delayed, unmeasured or normatively underdetermined.

5. Main Theorem

Theorem 1 - Fiber criterion for guaranteed correctness

Let $z \in Z$ and let $F_z = C^{-1}(z)$ be non-empty. There exists an availability-respecting deterministic decision rule $f: Z \rightarrow D$ that is correct for every state $x \in F_z$ at representation z if and only if:

$$A^i(z) \cap \bigcap_{x \in F_z} B(x) \neq \emptyset.$$

Equivalently, $G(z) \neq \emptyset$.

Proof. First, suppose that there exists an availability-respecting deterministic decision rule f that is correct for every $x \in F_z$. Since the rule operates on z , it selects a single decision $f(z)$. Because it is availability-respecting, $f(z) \in A^i(z)$. Because it is correct for every $x \in F_z$, $f(z) \in B(x)$ for every $x \in F_z$. Therefore $f(z) \in A^i(z) \cap \bigcap_{x \in F_z} B(x)$, so the intersection is non-empty.

Conversely, suppose that $A^i(z) \cap \bigcap_{x \in F_z} B(x) \neq \emptyset$. Then there exists a decision $d \in D$ such that $d \in A^i(z)$ and $d \in B(x)$ for every $x \in F_z$. Define a deterministic rule f such that $f(z) = d$. Then f is availability-respecting at z and correct for every state in F_z . This proves the equivalence. $\square$

Interpretation. The theorem gives the exact condition under which a representation z supports a guaranteed correct decision across all real states compatible with it. The condition is not merely that some correct decision exists for each state. The same available decision must be correct for all states in the fiber.

6. Structural Inaccessibility under Pairwise Conflict

Corollary 1 - Pairwise conflict implies no guaranteed correctness

Let $x_1, x_2 \in X$. Suppose $C(x_1) = C(x_2) = z$ and $B(x_1) \cap B(x_2) = \emptyset$. Then no deterministic decision rule $f: Z \rightarrow D$ can be correct for both x_1 and x_2 .

Proof. Since $C(x_1) = C(x_2) = z$, any deterministic rule selects the same decision for both states: $f(C(x_1)) = f(C(x_2)) = f(z)$. If f were correct for both states, then $f(z) \in B(x_1)$ and $f(z) \in B(x_2)$, hence $f(z) \in B(x_1) \cap B(x_2)$. This contradicts the hypothesis. $\square$

This corollary does not imply that $A^i(z) \cap B(x) = \emptyset$ for one of the two states. That stronger claim would be false in general. It may be that $A^i(z)$ contains a correct decision for x_1 and also contains a different correct decision for x_2 . The impossibility concerns guaranteed selection from the shared representation z , not the mere existence of state-specific correct actions.

7. State-Level Necessary Failure

Proposition 1 - Necessary failure is an intersection condition

For a state $x \in X$, necessary failure under the given representation and availability constraints occurs if and only if $A^i(C(x)) \cap B(x) = \emptyset$.

Proof. By definition, $A^i(C(x))$ is the set of decisions available under the representation of x . $B(x)$ is the set of decisions correct for x . Their intersection is the set of decisions that are both available and correct. If the intersection is empty, no available decision is correct. If it is non-empty, at least one available correct decision exists. $\square$

8. Corollaries and Consequences

If $G(z) = A^i(z) \cap \bigcap_{x \in F_z} B(x) = \emptyset$, then replacing a deterministic decision rule f with another deterministic rule g cannot produce guaranteed correctness across the fiber F_z , as long as C , A^i , and B remain unchanged.

If $G(z) = \emptyset$, guaranteed correctness across the fiber can be restored only by changing C so that the relevant states are separated, expanding or modifying $A^i(z)$, revising the task or correctness criterion B , introducing additional information not contained in z , or accepting that no guaranteed decision is available and shifting to risk management.

Aggregate accuracy does not identify whether errors are concentrated in decisionally incoherent fibers. Two systems may have identical accuracy but different structural risk. A system whose errors occur in fibers with empty $G(z)$ is limited by representation. A system whose errors occur despite non-empty $G(z)$ may be improved by better selection.

9. Probabilistic Extension

Let $P(x | z)$ be a posterior distribution over real states compatible with representation z . Let $L(d, x)$ be a loss function. A Bayes decision may be defined as:

$$d^i(z) \in \arg \min_{d \in A^i(z)} E[L(d, X) | z].$$

This decision may be optimal in expected loss. However, expected optimality does not imply guaranteed correctness for every state in the support of $P(\cdot | z)$.

If there exist x_1, x_2 in the support of $P(\cdot | z)$ such that $B(x_1) \cap B(x_2) = \emptyset$, then no deterministic decision based only on z can be correct for both states. The probabilistic rule can manage risk, distribute error, or minimize expected loss. It cannot recover a distinction absent from z .

10. Cumulative Representational Drift

Decision systems often involve chains of transformations rather than a single mapping. A clinical pathway converts symptoms into measurements, measurements into scores, scores into protocols. A legal process converts events into evidence, evidence into admissible records,

records into judicial findings. An AI pipeline converts inputs into features, features into embeddings, embeddings into outputs.

Let $X_0 \xrightarrow{C_1} X_1 \xrightarrow{C_2} X_2 \xrightarrow{C_3} \dots \xrightarrow{C_n} X_n$.

Proposition 2 - Irrecoverability of collapsed distinctions under deterministic downstream transformations

Suppose that for some k , two states $u, v \in X_{k-1}$ satisfy $C_k(u) = C_k(v)$. If all subsequent transformations $C_{k+1}, \dots, C_n$ operate only on the output of C_k , and no external information reintroduces a distinction between u and v , then the distinction between u and v cannot be recovered downstream.

Proof. After transformation C_k , both states have the same image. Every subsequent deterministic transformation receives identical input for both branches. A deterministic function maps identical inputs to identical outputs. Therefore all later representations remain identical unless information external to the transformation chain is introduced. $\square$

11. Representational Collapse Index

Let $S \subset X$ be a finite sample. The representation C partitions S into empirical fibers $E_S = \{e_1, e_2, \dots, e_m\}$. Each e_j contains sample cases with the same representation.

An empirical class $e \in E_S$ is incoherent if $\bigcap_{x \in e} B(x) = \emptyset$. A sufficient practical test is pairwise conflict: there exist $x_i, x_j \in e$ such that $B(x_i) \cap B(x_j) = \emptyset$.

The Representational Collapse Index is:

$$RCI(S) = \frac{|\{x \in S : x \text{ belongs to an empirically incoherent class}\}|}{|S|}.$$

RCI estimates the proportion of audited cases located in representational classes where no common correct decision exists. $RCI(S) = 0$ means that no incoherent collapse was detected in the sample under the chosen representation and correctness criterion. It does not prove absence of collapse outside the sample.

In metric or embedding spaces, exact equality may be too strict. Let ρ be a distance on Z , and let $\varepsilon > 0$. Define approximate equivalence by $\rho(C(x_i), C(x_j)) \leq \varepsilon$. The corresponding RCI_ε must always be reported with ε .

For a region $R \subseteq Z$, define a local index:

$$RCI_R(S) = \frac{|\{x \in S : C(x) \in R \text{ and } x \text{ belongs to an incoherent class}\}|}{|\{x \in S : C(x) \in R\}|}.$$

When a severity function $w(x) \geq 0$ is available, define:

$$RCI_w(S) = \frac{\sum_{x \in S} w(x) 1_{\{x \text{ belongs to an incoherent class}\}}}{\sum_{x \in S} w(x)}.$$

12. Audit Protocol

The framework can be applied as an audit procedure.

1. Define the real-state domain X .
2. Define the representation Z .
3. Specify the representation map C .
4. Define the available decisions $A^i(z)$.
5. Define the correctness correspondence $B(x)$.
6. Test state-level intersections $A^i(C(x)) \cap B(x)$.
7. Test fiber-level intersections $A^i(z) \cap \bigcap_{x \in F_z} B(x)$.
8. Classify cases as avoidable error, state-level necessary failure, fiber-level structural ambiguity, or indeterminate.
9. Estimate RCI, local RCI, and any severity-weighted extension.
10. Locate intervention in representation, availability, correctness criteria, data acquisition, procedure, or downstream risk management.

13. Controlled Computational Illustration

Let the real state include an observed component u and a hidden component h : $x = (u, h)$. Let the representation discard the hidden component: $C(u, h) = u$. Let the correct decision depend on h : $B(u, 0) = \{d_0\}$, $B(u, 1) = \{d_1\}$, and $d_0 \neq d_1$.

Then $C(u, 0) = C(u, 1) = u$, but $B(u, 0) \cap B(u, 1) = \emptyset$. By Corollary 1, no deterministic decision rule based only on u can be correct for both states. This illustrates structural inaccessibility in its simplest form.

In an empirical notebook, this structure can be tested by generating or identifying cases where reduced representations collide while labels or correct decisions conflict. Any numerical result must be reported as dataset-specific, not as proof of a universal law.

14. Domain Applications

14.1 Clinical decision-making

In clinical decision-making, X is the patient's real physiological state. Z includes observed measurements, clinical signs, diagnostic images, laboratory values and medical records. The map C includes measurement, timing, clinical attention, available technology and documentation.

Two patients may share the same representation under available tests while differing in an unobserved variable relevant to treatment. If the correct treatment differs between those

patients, the clinician faces a representation-induced limit. The framework does not imply that clinical responsibility disappears. It distinguishes the level at which responsibility should be examined.

14.2 Legal decision-making

In legal systems, X is the full historical event or dispute. Z is the procedural representation: admitted evidence, filed documents, pleadings, judicially recognized facts and procedural constraints. The map C includes filing rules, evidentiary admissibility, deadlines, translation, record preservation and institutional selection.

The legal file is not the full reality. It is a constructed representation. If two materially different histories become indistinguishable in the file while requiring different legal outcomes, the court may be unable to guarantee a correct decision from the file alone. This does not automatically prove judicial error. It identifies where the risk arises: in the construction of the procedural representation.

14.3 Artificial intelligence

In artificial intelligence, X is the underlying task-relevant state. Z may be a feature vector, embedding, token sequence, latent representation or sensor output. The map C may be manual feature extraction, dimensionality reduction, learned encoding, compression, tokenization or data preprocessing.

If different states with different correct labels map to the same representation, downstream model improvement cannot guarantee separation. A more complex classifier cannot recover a distinction not present in its input. This is especially relevant for fairness, safety and out-of-distribution failures.

15. Discussion

The framework separates four questions. Was the decision wrong? Was a correct decision available under the representation? Could a single decision be guaranteed correct for all states compatible with the representation? Did the representation preserve the distinctions required by the correctness criterion?

Many audits stop at the first question. This is insufficient. A wrong decision may be an avoidable selection error. It may also be the manifestation of a deeper structural limit.

The framework also explains why optimization can plateau. If errors are caused by decisionally incoherent fibers, changing the downstream rule can redistribute errors but cannot eliminate them. The correct intervention is upstream: improve measurement, preserve evidence, add features, change representation, expand available actions, or redefine the decision architecture.

16. Limitations

First, $B(x)$ is domain-dependent. Correctness in medicine, law and AI is not defined by the same normative structure. Second, X is rarely fully observable. Real states may be reconstructed imperfectly or only after the fact. Third, C may be unknown or only partially inspectable. In learned systems, representation may be distributed and opaque.

Fourth, the main theorem concerns deterministic decision rules. Randomized rules can alter risk distributions, but they do not create a guaranteed correct decision when no common available correct decision exists across a fiber. Fifth, RCI is sample-dependent. A low RCI does not prove the absence of structural collapse outside the sample. Sixth, the framework does not assign moral, legal or institutional responsibility. It identifies structural conditions that may be relevant to responsibility analysis.

The paper does not claim a universal closed mathematical theory of decision correctness. A fully general theory would require domain-independent definitions of states, correctness, observability, measurement and normativity. Such definitions are not available for heterogeneous real systems.

17. Conclusion

This paper formalized decision accessibility under representation constraints. The central result is a fiber criterion: a guaranteed correct decision exists for all real states compatible with a representation z if and only if the available decision set intersects the common correctness set of the fiber $C^{-1}(z)$.

When this intersection is empty, no deterministic rule operating only on z can guarantee correctness across the fiber. In pairwise form, if two states collapse into the same representation while requiring mutually incompatible correct decisions, at least one of them must be decided incorrectly by any deterministic rule based only on that representation.

The framework distinguishes avoidable error, state-level necessary failure, fiber-level structural ambiguity and indeterminate cases. It also introduces RCI as an operational audit metric for detecting representational collapse in finite systems.

The main implication is practical: when failure is structural, better choice within the same representation is not enough. The representation itself must be examined.

18. Preprint Status and Reproducibility Statement

This document is intended for archival deposit as a working paper/preprint. The framework is theoretical and does not rely on unpublished empirical data. Version 1.0 should be treated as a citable archival version. Any later correction or extension should be deposited as a new Zenodo version rather than silently replacing the scientific record.

Any future computational notebook, dataset-specific audit or domain application should be deposited as a separate supplementary file and should specify the finite sample used, the representation map or observed representation, the operational definition of correctness, the availability constraints, the equivalence or similarity threshold, and the resulting RCI, local RCI and any severity-weighted extension.

19. Ethics, Responsibility and Domain Boundaries

The framework can be applied to clinical, legal, administrative and artificial intelligence systems, but it does not replace domain-specific standards of care, legal standards of proof, judicial reasoning, medical diagnosis, or regulatory compliance. It is an analytical structure for

identifying whether a failure concerns selection within a representation or the adequacy of the representation itself. In high-stakes domains, the framework should be used as an audit aid, not as an automatic decision tool.

20. Data Availability

No empirical dataset is used in this theoretical working paper. The examples are formal and illustrative. Any future empirical implementation should provide the dataset, code, assumptions and thresholds necessary to reproduce the reported RCI values.

21. Conflict of Interest Statement

The author declares no institutional conflict of interest in relation to this working paper.

22. Funding Statement

This work received no external funding.

23. Suggested Zenodo Metadata and Deposit Checklist

Title: Decision Accessibility and Structural Inaccessibility under Representational Collapse

Creator: Pierernesto Di Paolo

ORCID: 0009-0003-0149-4321

Resource type: Preprint / Working paper

Language: English

Keywords: decision theory; representational collapse; structural inaccessibility; bounded rationality; information loss; audit; artificial intelligence; legal decision-making; clinical decision-making; RCI

License: CC BY 4.0 is recommended for maximum scientific reusability and is the standard open-access choice for a preprint; a more restrictive custom license may be used only if intentional.

Version: 1.0

DOI strategy: use the version-specific DOI for citing this exact version; use the concept DOI only when referring to all future versions collectively.

Before publication on Zenodo, verify that the title is identical in the PDF and Zenodo metadata; the author name and ORCID are correct; the resource type is set to preprint or working paper; the license is intentionally selected; the version is marked as 1.0; no placeholder DOI remains inside the PDF; the PDF is final and internally consistent; supplementary files are clearly named and versioned; the description field reproduces the abstract or a shortened version of it; and the keywords match the paper's actual content.

References

Arrow, K. J. (1951). *Social Choice and Individual Values*. New York: Wiley.

Ashby, W. R. (1956). *An Introduction to Cybernetics*. London: Chapman & Hall.

Bellman, R. (1957). *Dynamic Programming*. Princeton: Princeton University Press.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2), 265-272.

Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Hoboken: Wiley.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379-423; 27(4), 623-656.

Simon, H. A. (1957). *Models of Man: Social and Rational*. New York: Wiley.

Wald, A. (1950). *Statistical Decision Functions*. New York: Wiley.

Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press.

Appendix A - Minimal Audit Pseudocode

Input: finite sample S ; representation map C or observed representation z ; correctness sets $B(x)$ or labels/proxy correctness; availability sets $A^i(z)$, if available; optional distance threshold ε .

Procedure: compute representations $z = C(x)$ for all x in S ; group cases by identical z or approximate equivalence; for each group, compute whether the common correctness intersection is empty; mark all cases in incoherent groups; compute RCI; if $A^i(z)$ is known, compute $G(z)$; classify groups as guaranteed accessible, structurally ambiguous, state-level necessary failure, or indeterminate.

Output: RCI; list of incoherent classes; local RCI by region; recommended intervention point.

Appendix B - Terminology

Real state X : the underlying condition about which a decision is made.

Representation Z : what the system actually observes or uses.

Representation map C : the transformation from real state to representation.

Available decisions $A^i(z)$: decisions available under representation z .

Correct decisions $B(x)$: decisions correct for real state x .

Fiber F_z : all states mapped to the same representation z .

Representational collapse: multiple real states mapped to the same representation.

Decisionally incoherent fiber: a fiber whose states have no common correct decision.

Avoidable error: wrong decision despite existence of a correct available decision.

Necessary failure: absence of any correct available decision for a state.

Fiber-level structural ambiguity: absence of a single available decision guaranteed correct for all states in a fiber.

RCI: finite-sample index measuring the proportion of cases located in incoherent representational classes.